arXiv:1610.10028v1 [math.ST] 31 Oct 2016

Confidence intervals with control of the sign error in low power settings

Art B. Owen
Stanford University

October 2016

Abstract

When hypothesis tests of $H_0 : \theta = 0$ have low power, it is possible that their rejection can frequently be accompanied by an estimate $\hat{\theta}$ that has the wrong sign and significantly exaggerated magnitude. Such sign errors are less likely when the confidence interval for θ is well separated from 0, as measured in units of the confidence interval's width. Sign errors can be controlled by declaring confidence in the sign of θ only when H_0 is also rejected at a smaller level $\alpha_2 = 2\alpha_1\alpha_S$, where $\alpha_S \le 1/2$ is a user specified upper bound on the probability of a sign error given that H_0 has been rejected. This procedure is a very simple form of inference after model selection: the selected model is that $\theta \ne 0$ and the second inference is then on sign(θ).

1 Introduction

A serious problem arises for hypothesis tests of low statistical power, as pointed out by Gelman (2014). Low power arises when sampling uncertainty is large compared to the effect size. Statistical significance requires the estimated effect size to be a reasonable multiple of the sampling uncertainty. Therefore all of the observed significant outcomes exaggerate the true effect size, in this low power setting.

The issue typically arises in the context of testing a null hypothesis $H_0 : \theta = 0$ for some parameter θ of interest. It is usually implausible that θ could be exactly zero in a real world context. A common justification for testing H_0 anyway is that when you reject it, at least you have good information about the sign of θ. If $\theta > 0$, then testing against $\theta = -\epsilon$ is a relatively hard negative value to distinguish from θ and H_0 emerges in the limit as $\epsilon \to 0$.

Gelman's example shows that in the low power case, rejecting the point null H_0 can not only exaggerate the magnitude of θ, it can also leave one with an incorrect sign. The point was made earlier by Gelman and Tuerlinckx (2000) who showed that some hierarchical Bayesian credible intervals greatly diminish the sign error problem, although their inferences are more conservative.

One way to handle the problem is to avoid doing tests that have low power. We should instead use a large sample or some other informative design and only study effects that are large. However we often cannot be sure what the power is going to be before doing a study and so low power settings will inevitably arise. Another approach is to employ Bayes factors, as for example Bayarri et al. (2016) do in a Bayes-frequentist hybrid. That will not suit every user as specifying priors can be difficult.

Confidence intervals have the benefit of being in the same units as the effect θ and hence they facilitate study of practical significance. The approach proposed here is to only draw conclusions about the sign of $\hat{\theta}$ when a confidence interval for θ is not very wide compared to the distance from $\hat{\theta}$ to zero. This method takes explicit account of the apparent magnitude of θ, and not just whether the interval includes zero. In low power settings, many rejections of H_0 will be unaccompanied by a decision about the sign of θ. In high throughput testing of many hypotheses, some but not all rejections may be accompanied by decisions about the sign.

The approach is an extremely basic form of conditional inference as described by Fithian et al. (2014) who also give a comprehensive bibliography. We reject H_0 at some level α_1 to conclude that $\theta \neq 0$ and if we can also reject H_0 at a more stringent level α_2, then we additionally conclude that $\text{sign}(\theta) = \text{sign}(\hat{\theta})$.

An outline of this paper is as follows. Section 2 gives the problem formulation and shows how to compute the standard error of $\hat{\theta}$ given the level α and power $1-\beta$ of a test of $H_0 : \theta = 0$ when the true θ is one unit. It seems that there is an error in the figure from Gelman (2014) but that his substantive points remain valid. After an email conversation with Andrew Gelman, it appears that figure illustrated a data set in which the power was just over 5.5% which then rounds up to 6%. Section 3 shows how sign errors and exaggeration factors vary with the underlying power when $\alpha = 0.05$. That is not to say 0.05 is a good choice. It is however, the popular one. Sign errors become very unlikely when the power is over 30% but magnitude errors can be substantial at higher powers. Section 4 proposes a way to control sign errors conditionally on rejecting H_0. We motivate it by comparing the confidence interval width to the distance of $\hat{\theta}$ from 0. The idea is to choose two signficance levels α_1 and $\alpha_S < 1/2$. If H_0 can be rejected at level α_1 conclude that $\theta \neq 0$. If H_0 can also be rejected at the level $\alpha_2 = 2\alpha_1\alpha_S$, conclude that we also know the sign of θ equals that of $\hat{\theta}$. The probability of a sign error given that H_0 has been rejected is then at most α_S. The power for conclusions that come with a decision on the sign of θ is necessarily lower than the power for simply rejecting H_0. This seems appropriate. It is also possible to define a p-value for the sign as the smallest α_S for which H_0 is rejected at level α_2, although this approach requires that α_1 be prespecified and not a researcher degree of freedom. Section 5 looks at one-tailed tests. They do not solve the problem. For the same effect size that has 6 percent power, running a one-tailed test at level $\alpha = 0.05$ would (falsely) reject H_0 just over 26% of the time.

```
tau = function(alpha,powr){
# A Gaussian test at level alpha
# has power powr for Y ~ N( 1, tau^2 ).
# Solve for tau using noncentral chisquare.

aux = function(tau){
  1-powr - pchisq(qchisq(1-alpha,1),1,ncp=1/tau^2)
}

ur = uniroot( aux, lower=10^-9,upper=10^6,
    tol = .Machine$double.eps^0.9)
ur$root
}
```

Figure 1: R code to compute the critical standard error τ.

2 Problem formulation

We focus on the large sample setting where the statistical estimate is approximately normally distributed. Suppose that $\theta \in \mathbb{R}$ is an effect that we estimate by a statistic $\hat{\theta}$. There is a null hypothesis H_0 that $\theta = \theta_0$. The most common case has $\theta_0 = 0$ which we assume from here on. We reject H_0 for a large $|\hat{\theta}|$.

Let $\hat{\theta} \sim \mathcal{N}(\theta, \tau^2)$ and suppose there is a variance estimate $s^2 \sim \tau^2\chi^2_{(\nu)}/\nu$ with a large number ν of degrees of freedom, independent of $\hat{\theta}$. We choose our units so that the effect size $\theta = 1$. We reject the null hypothesis $\theta = 0$ at level $\alpha \in (0, 1)$ if

$$\frac{\hat{\theta}^2}{s^2} = \frac{(\hat{\theta}/\tau)^2}{(s/\tau)^2}$$

exceeds the $1-\alpha$ quantile of the F distribution with 1 numerator and ν denominator degrees of freedom. Since ν is large, the test asymptotically rejects H_0 when $\hat{\theta}/\tau$ exceeds the $1-\alpha$ quantile of the $\chi^2_{(1)}$ distribution.

The independence of s^2 from $\hat{\theta}$ won't always hold. However, we get the same result from $\hat{\theta}$ being asymptotically $\mathcal{N}(\theta, \tau^2)$ and s^2 converging in distribution to τ^2.

To study the alternative hypothesis $\theta \neq 0$ it is convenient to scale the measurement units of the problem so that $\theta = 1$. Then $Z \equiv \hat{\theta}/\tau \sim \mathcal{N}(1/\tau, 1)$ and we reject H_0 if $Z^2 \sim \chi'^2_{(1)}(\tau^{-2}) \geq \chi^{2,1-\alpha}_{(1)}$ where χ'^2 denotes a noncentral χ^2 distribution.

For a given significance level $\alpha \in (0, 1)$ and power $1-\beta \in (\alpha, 1)$, we can solve for the standard error τ. Some R code to do this is in Figure 1. For $\alpha = 0.05$ and $1-\beta = 0.06$ we find that $\tau = 3.394507$. The point of all those digits in τ is that we can plug this value in to the `retrodesign` R function from Gelman and Carlin (2014) and verify that it yields power 0.06 when $\alpha = 0.05$. See Figure 2.

```
> retrodesign
function(A, s, alpha=.05, df=Inf, n.sims=10000){
z <- qt(1-alpha/2, df)
p.hi <- 1 - pt(z-A/s, df)
p.lo <- pt(-z-A/s, df)
power <- p.hi + p.lo
typeS <- p.lo/power
estimate <- A + s*rt(n.sims,df)
significant <- abs(estimate) > s*z
exaggeration <- mean(abs(estimate)[significant])/A
return(list(power=power, typeS=typeS, exaggeration=exaggeration))
}

> set.seed(1);retrodesign(1,3.394507)
$power
[1] 0.06

$typeS
[1] 0.2013426

$exaggeration
[1] 7.978919
```

Figure 2: R code from Gelman and Carlin (2014) to verify that standard error $\tau = 3.394507$ and effect size $\theta = 1$ yields power 0.06 at $\alpha = 0.05$. The random seed is set to make the results more reproducible.

The figure in Gelman (2014) (as of October 22, 2016) did not match these results. It had a wrong sign probability of 24% (versus about 20% in Figure 2) and a minimum exaggeration factor of 9 (versus an average of about 8). It appears that the figure in the blog is based on slightly different numbers.

For $\alpha = 0.05$ and power 0.06, using τ from the code in Figure 1 leads to a critical value of $\Phi^{-1}(0.975) \times \tau \doteq 6.65$. Therefore any significant discovery must overestimate θ by at least 6.65-fold. The average value of $|\hat{\theta}/\theta|$ given $\alpha = 0.05$ and power 0.06 is 8.01 and the probability of a sign error (given rejection of H_0) is about 20%. These values closely match the ones from Figure 2 making it more likely that the blog figure is for some power level close to but not exactly 0.06.

3 Error versus power

Here we consider the prevalance of sign errors when $\alpha = 0.05$. We vary the power from the low level of 0.06 to the high level 0.80 that is often used as a design goal. Figure 4 shows the conditional probability of a sign error given

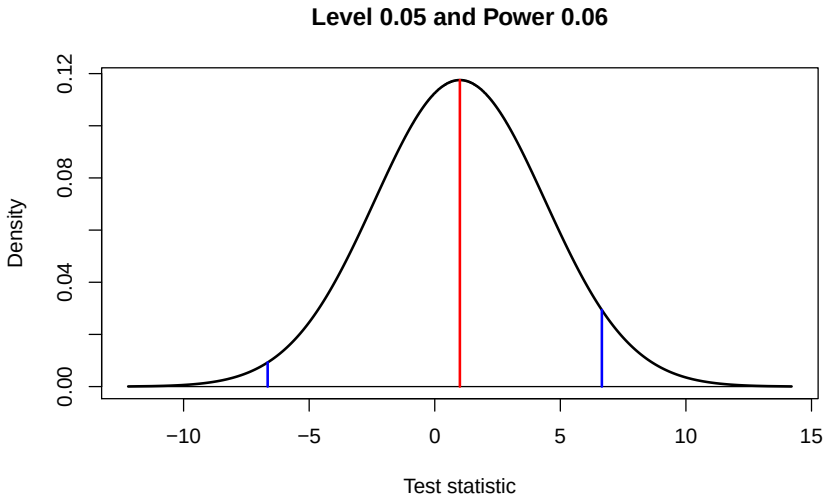

Figure 3: The curve shows the $\mathcal{N}(1, \tau^2)$ density of $\hat{\theta}$ when a test of $H_0 : \theta = 0$ at level 0.05 has power 0.06. The central vertical line is at the true value $\theta = 1$. The other vertical lines are at $\hat{\theta} = \pm 6.65$. H_0 is only rejected when $|\hat{\theta}| \geq 6.65$. About 20% of rejections have the wrong sign.

that H_0 is rejected. It becomes negligible at power greater than about 30%.

Next we consider magnitude errors in $\hat{\theta}$. The light curve in Figure 5 shows the minimal value of $|\hat{\theta}/\theta|$ for which H_0 can be rejected at level 0.05. Whenever the power is below 0.5 this minimal ratio is above 1. When the power exceeds 50% then H_0 can be rejected with either an exaggerated magnitude or an underestimated magnitude. The darker curve in Figure 5 shows the expected value of $|\hat{\theta}/\theta|$ given that H_0 has been rejected, versus the power of the test. Even at 50% power the average exaggeration can be as high as 1.4 fold.

4 Confidence in the sign

If the confidence interval for θ just barely excludes zero then we should be more cautious about the sign of θ than if the distance from θ to 0 is a large multiple of the confidence interval width.

To illustrate, suppose that we reject H_0 when $|\hat{\theta}| \geq \Phi^{-1}(0.975)\tau \doteq 1.96\tau$. Suppose further that we only declare the sign of θ to match that of $\hat{\theta}$ when the center $\hat{\theta}$ of the confidence interval is at least $2 \times 1.96\tau$ away from 0. That is, we require the separation between the confidence interval and 0 to be at least half of the confidence interval width. Under H_0 such separation only happens with probability $2\Phi(-3.92) \doteq 8.85 \times 10^{-5}$. Our chance of wrongly declaring the sign

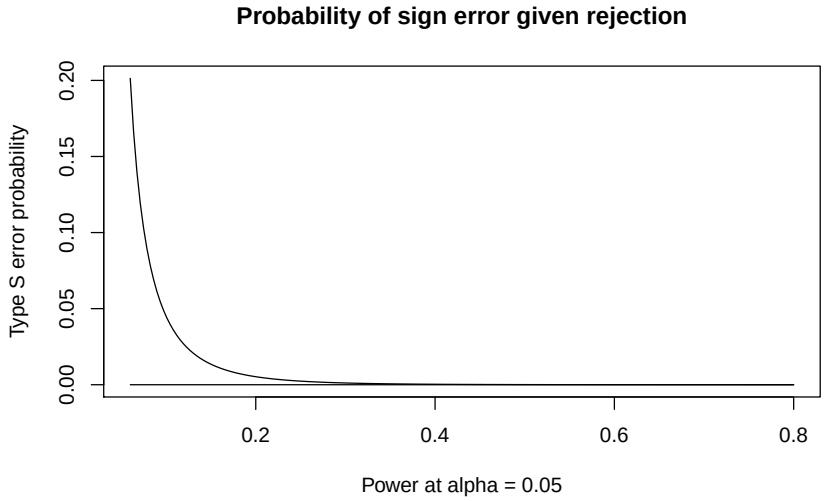

Figure 4: The figure shows how the probability of a sign error, conditional on rejecting H_0, depends on power.

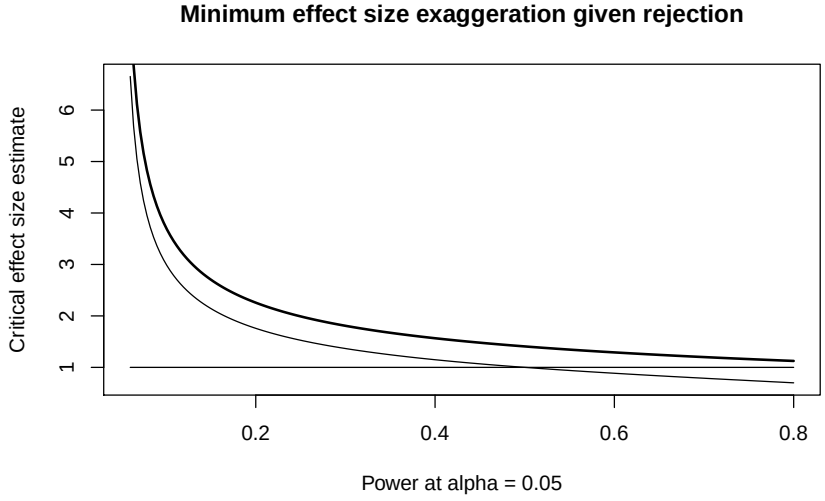

Figure 5: The horizontal axis is the power of a test at level $\alpha = 0.05$. The solid curve is the expected value of $|\hat{\theta}/\theta|$ given that H_0 has been rejected. The lighter curve is the minimum value of $|\hat{\theta}/\theta|$ for which H_0 could be rejected.

given that we have rejected H_0 is now at most

$$8.85 \times 10^{-5}/0.05 \doteq 0.0018.$$

This estimate 0.0018 assumes that every time we reject at $\alpha = 0.05$ but fail to reject at $\alpha' = 2\Phi^{-1}(-3.92)$ implies a sign error. The fraction of sign errors is at most 0.5 in the Gaussian setting and is only about 20% when $\alpha = .05$ and the power is 0.06. Therefore a more realistic, yet still conservative estimate of the conditional probability of a sign error given that we have rejected H_0 and the confidence interval endpoint is at $\hat{\theta}/2$ is $0.5 \times 8.85 \times 10^{-5}/0.05 \doteq 0.0009$, just under one in 1000.

Now suppose that H_0 is rejected at primary significance level α_1. Next we stipulate a secondary significance level $\alpha_2 < \alpha_1$. Let R_1 be a binary variables taking the value 1 iff we reject H_0 at α_1 and let R_2 be the binary variable for α_2. The probability that we make a sign error under this rule given that we have rejected H_0 is

$$\frac{\Pr(\mathrm{sign}(\hat{\theta}) \neq \mathrm{sign}(\theta))}{\Pr(R_1 = 1)} \leq \frac{1}{2}\frac{\Pr(R_2 = 1)}{\Pr(R_1 = 1)} = \frac{1}{2}\frac{\alpha_2}{\alpha_1}.$$

Given a target conditional sign error of α_S we use $\alpha_2 = 2\alpha_1\alpha_S$. For instance, if we want a type I error rate of $\alpha_1 = 0.01$ and a conditional sign error rate of at most 0.01, then we take $\alpha_2 = 2 \times 10^{-4} = 1/5000$.

Suppose for example that we want to test H_0 at level $\alpha = 0.05$ but we are very averse to sign errors and want $\alpha_S = 0.001$. That is at most one in one thousand rejections of H_0 should come with a sign wrongly declared by the test at level $\alpha_2 = 0.0001$. Figure 6 shows both the power to reject H_0 and the power to identify $\mathrm{sign}(\theta)$ as a function of $|\theta|/\tau$. It also depicts more lenient sign conditions given by $\alpha_S = 0.01$ and $\alpha_S = 0.1$. Ordinarily $\tau \doteq \sigma/\sqrt{n}$ where σ is an asymptotic standard deviation. There is a range of effect sizes in which rejections of H_0 without decisions on $\mathrm{sign}(\theta)$ will be common. But once $|\theta|\sqrt{n} \geq 5\sigma$ or so, there is very high power that H_0 will be rejected and a sign will be identified. Also, the hypotheses in limbo with H_0 rejected but no sign determination are the ones for which obtaining additional data may be most valuable.

In high throughput settings, we might make a very large number of primary hypothesis tests at level α_1. For each one that is rejected, there is then a smallest $\alpha_2 \leq \alpha_1$ at which H_0 is also rejected. Call this value p_2. Then the smallest α_S for which we would have found the sign significant is $p_S = p_2/(2\alpha_1)$. Note that p_S depends on the chosen value α_1, which should then be regarded as a fixed pre-specified choice, and not as a function of the observed data.

5 One-sided tests

Here we investigate the possibility of sign errors in one-tailed tests. Those tests are sometimes justified by an assumption that the direction of the effect

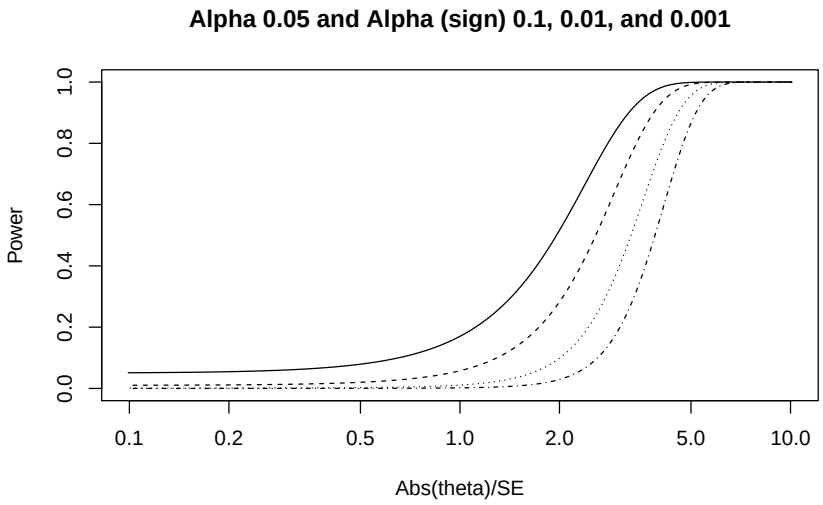

Figure 6: The horizontal axis is the effect size $|\theta|/\tau$. The solid curve is the power of a test of $H_0 : \theta = 0$ at $\alpha = 0.05$. The next curves from top to bottom are the probabilities of also making a definitive sign declaration at $\alpha_S \in \{0.1, 0.01, 0.001\}$.

is certain a priori. Then sign errors are not possible, unless our certain opinion is wrong. The a priori uncertainty about the sign of θ must however be small compared to the critical level α in use, and that is a very strong assumption to work under.

A second justification is that sometimes only one direction is consequential. An opinion about one direction being inconsequential might also be mistaken, but for sake of argument we work with it. In this situation we might well make a sign error by rejecting H_0. Suppose that $\hat{\theta} \sim \mathcal{N}(1, \tau^2)$, so that $\hat{\theta}/\tau \sim \mathcal{N}(1/\tau, 1)$. The true effect is positive but we might be testing for a (consequential) negative effect.

A one-tailed test at level α in the negative direction would reject H_0 if $\hat{\theta}/\tau \le \Phi^{-1}(\alpha)$. This happens with probability $\Phi(\Phi^{-1}(\alpha) - 1/\tau)$.

Now we revisit the case of 6% power for a two-tailed test at $\alpha = 0.05$. Then $\tau \doteq 3.39$ and the sign error is 0.20. The wrong direction one-tailed test at level 0.05 will reject H_0 with probability $\Phi(\Phi^{-1}(0.05) - 1/\tau) \doteq 0.026$.

The chance of a wrong sign rejection is actually larger than $\alpha/2$ here. That is not to say the conditional probability of a rejection being wrong is over 50%. Indeed in this setting any rejection at all is a sign error. What it does mean is that in this low power setting, the probability of a sign error is not small compared to α. As a result, using one-tailed tests does not correct the problem of sign errors in low power settings.

Acknowledgments

I thank Jelena Markovic and Andrew Gelman for discussions. This work was supported by the NSF under grants DMS-1407397 and DMS-1521145.

References

Bayarri, M. J., Benjamin, D. J., Berger, J. O., and Sellke, T. M. (2016). Rejection odds and rejection ratios: A proposal for statistical practice in testing hypotheses. *Journal of Mathematical Psychology*, 72:90–103.

Fithian, W., Sun, D., and Taylor, J. (2014). Optimal inference after model selection. Technical report, arXiv:1410.2597.

Gelman, A. (2014). *This is what "power = .06" looks like. Get used to it.* `http://andrewgelman.com/2014/11/17/power-06-looks-like-get-used/`.

Gelman, A. and Carlin, J. (2014). Beyond power calculations assessing type S (sign) and type M (magnitude) errors. *Perspectives on Psychological Science*, 9(6):641–651.

Gelman, A. and Tuerlinckx, F. (2000). Type S error rates for classical and Bayesian single and multiple comparison procedures. *Computational Statistics*, 15(3):373–390.

Additional R code

The code here was used to compute the expected value of $\hat{\theta}$ conditionally on H_0 being rejected and on sign($\hat{\theta}$).

```
exag = function(alpha=.05,powr=.06){
# Compute exaggeration factors

critse = tau(alpha,powr)  # Critical Standard Error
critz  = critse*qnorm(1-alpha/2)
posmean = gauscondmean(mu=1,sigma=critse,A=critz, B=Inf, n=10^5 )
negmean = gauscondmean(mu=1,sigma=critse,B=-critz, n=10^5 )
typeS   = pnorm(-critz,1,critse)/powr

exager = posmean*(1-typeS)+abs(negmean)*typeS

list(critse=critse,critz=critz,exager=exager,posmean=posmean,negmean=negmean,typeS=typeS)
}


stdgauscondmean = function(A=-Inf, B=Inf, n=10^5, plot=FALSE ){
# Average of N( 0, 1 ) over interval (A,B)

u = ((1:n)-0.5)/n
u = pnorm(A) + (pnorm(B)-pnorm(A)) * u
z = qnorm(u)
if( plot )hist(z,50) # for testing/debug
mean(z)
}

gauscondmean = function(mu=0,sigma=1,A=-Inf, B=Inf, n=10^5 ){
# Average of N( mu, sigma^2 ) over interval (A,B)

mu + abs(sigma)*stdgauscondmean((A-mu)/abs(sigma),(B-mu)/abs(sigma),n)
}
```